



DualSpeech: Enhancing Speaker-Fidelity and Text-Intelligibility Through Dual Classifier-Free Guidance

Jinhyeok Yang^{1*}, Junhyeok Lee^{1*}, Hyeong-Seok Choi^{2†}, Seunghoon Ji¹, Hyeongju Kim¹, Juheon Lee¹

¹Supertone Inc., Republic of Korea

²ElevenLabs Inc., USA

yangyangii@supertone.ai, jun.hyeok@supertone.ai

Abstract

Text-to-Speech (TTS) models have advanced significantly, aiming to accurately replicate human speech's diversity, including unique speaker identities and linguistic nuances. Despite these advancements, achieving an optimal balance between speaker-fidelity and text-intelligibility remains a challenge, particularly when diverse control demands are considered. Addressing this, we introduce DualSpeech, a TTS model that integrates phoneme-level latent diffusion with dual classifier-free guidance. This approach enables exceptional control over speaker-fidelity and text-intelligibility. Experimental results demonstrate that by utilizing the sophisticated control, DualSpeech surpasses existing state-of-the-art TTS models in performance. Demos are available at <https://bit.ly/48Ewoib>.

Index Terms: text-to-speech, diffusion, classifier-free guidance

1. Introduction

Human speech is characterized by a wide array of variations, including distinctive speaker identities, different speech rhythms, tones, languages, and more. The goal of Text-to-Speech (TTS) is to emulate this richness, synthesizing speech that is not only natural and human-like but also encompasses a broad spectrum of these qualities and nuances.

Accordingly, state-of-the-art TTS models should excel in producing speech that not only captures the essence of the speaker, including their timbre, speaking style, accent, and emotions, for high speaker-fidelity but also ensures that the speech is easily understood, maintaining strong text-intelligibility [1, 2, 3, 4, 5, 6]. However, achieving a perfect balance between speaker-fidelity and text-intelligibility can be challenging in some cases. For example, using a yawning young woman's recording as a reference for speech synthesis might lead us into a dilemma: focusing too much on matching her voice, including the yawn, might compromise the clarity of the speech (text-intelligibility). Conversely, concentrating on making the speech clear could result in losing the yawn's unique effect, thereby producing speech that accurately represents a young woman's voice but lacks the intended distinctive characteristic, thus affecting speaker-fidelity.

Most TTS research to date evaluates these two components—speaker-fidelity and text-intelligibility—using metrics such as speaker similarity, naturalness MOS, and word error rate (WER); however, there has been little exploration

into methods for independently controlling each element when they come into conflict. We believe that the ability to independently manipulate these factors would be highly beneficial in real-world TTS scenarios. To address this, we have sought out methodologies that enable such control, focusing on: 1) the use of representation disentanglement within generative models to separate and independently manage different aspects of speech, and 2) the application of classifier-free guidance in diffusion-based generative models, which allows for the independent conditioning and control of various conditions, actively exploring these approaches for practical solutions.

First, NANSY¹ [7, 8] has demonstrated quality improvements through self-supervised reconstruction from disentangled features. Notably, NANSY stands out for providing high controllability via interpretable features such as linguistic features, fundamental frequency (f_0), periodic and aperiodic amplitudes, and timbre features. In particular, NANSY-TTS, an application of NANSY for TTS tasks, exemplifies NANSY's ability to independently manage timbre features, allowing for the disentanglement of speaking style and timbre. This capability affords enhanced controllability over the representation of various speakers. However, similar to broader challenges in the field, NANSY-TTS still grapples with controlling the balance between speaker-fidelity and text-intelligibility.

Second, diffusion models [9] employ classifier-free guidance (CFG) [10] to control various conditions, a technique also adopted by speech diffusion models [2, 11] for enhanced condition manipulation. Specifically, VoiceLDM [11] introduces dual CFG mechanism that allows separate control of environmental and content conditions. This feature allows VoiceLDM to manipulate the intensity of environmental and content conditions independently. While our approach is similar to dual CFG of VoiceLDM, our method is more ideal in TTS by enabling control between text-intelligibility that follows text content and speaker-fidelity of reference speech.

In this paper, we introduce DualSpeech, a latent diffusion-based TTS model that achieves enhanced speaker-fidelity and text-intelligibility by utilizing dual classifier-free guidance. To attain high controllability with dual CFG, we introduce two phoneme-level conditioners; reference conditioner and text conditioner. These networks are designed to model prior latent highly dependent on reference speech and text, respectively. Through these networks, at the inference stage, we can manipulate the prosody of the generated speech to follow either the reference or the content by selecting CFG weights. Our proposed approach demonstrates superior zero-shot TTS capability, along with enhanced intelligibility and controllability.

¹For clarity, all NANSY referred to in this paper is NANSY++ [7], rather than its earlier version [8].

*Equal contribution †Work done at Supertone Inc.

This research was supported by Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2022 (Project Name: Cultural Technology Research and Development, Project Number: R2022020066, Contribution Rate: 50%)

2. Method

DualSpeech is composed of three main components: NANSY, variational auto-encoder (VAE) [12], and latent diffusion model (LDM) [13]. The comprehensive architecture is illustrated in Figure 1. Unless specified otherwise, this paper assumes that the referenced modules are Transformer encoders, for which architectural details have not been provided.

Different from almost of TTS model, which generates mel-spectrogram, DualSpeech utilizes NANSY features. We leverage a pre-trained NANSY for extracting NANSY features, aligning with the NANSY-TTS, which generates linguistic feature, f0, periodic amplitude, and aperiodic amplitude [7]. In addition, Aligner is trained to align NANSY linguistic features with phonemes using monotonic alignment search (MAS) [14]. Building on these pre-trained models, our model is trained in two stages, VAE training, and LDM training. Our VAE, featuring a phoneme-level bottleneck, reconstructs NANSY features from given speech and phonemes. Lastly, LDM generates VAE latent from given transcription and reference speech.

2.1. Phoneme-Level Variational Auto-Encoder

The VAE in DualSpeech processes inputs comprising an IPA sequence, which has been converted from text, and the corresponding NANSY features of the speech, to reconstruct the NANSY features of that speech. Our VAE utilizes a phoneme-level bottleneck, inspired by [15, 16]. This bottleneck is implemented through the cross-attention mechanism of the Transformer encoder, which uses the output of the phoneme encoder as a query and concatenated NANSY features as both key and value. This approach offers two advantages over frame-level models. Firstly, a phoneme-based representation reliably conveys semantic information, as phonemes are symbolic representations of speech sounds. Secondly, it provides computational efficiency compared to models that learn at the frame-level, as the computation complexity of the Transformer encoder scales as $\mathcal{O}(L^2)$ [17], where L is the sequence length. From the phoneme-level bottleneck, the posterior latent is sampled by estimated mean and variance.

The VAE decoder comprises a latent decoder, a duration predictor, a phoneme prosody decoder, an upsampler, and a frame decoder. The decoding process begins with the latent variables passing through the latent decoder. The outputs from this network feed into the duration predictor, phoneme prosody decoder, and upsampler in parallel. The duration predictor and phoneme prosody decoder are responsible for estimating duration and f0 at the phoneme-level, respectively. At the upsampler stage, phoneme-level output of the latent decoder is upsampled to a frame-level sequence by ground truth duration from the pre-trained MAS aligner. The architecture of the upsampler is almost identical to the learned upsampler from Parallel Tacotron 2 [18], excluding the channel dimension. The upsampled frame-level feature is then input into the frame decoder to reconstruct NANSY features.

In addition, we enhance the performance of the VAE through adversarial training [19, 20]. Our discriminator consists of a simple convolution network trained using least-square loss and feature-matching loss.

Our VAE model is trained with NANSY feature reconstruction losses, phoneme-level f0 reconstruction loss, duration loss, KL divergence of latent, and adversarial losses.

2.2. Phoneme-Level Latent Diffusion Model

DualSpeech’s LDM is trained to estimate the phoneme-level posterior latent generated by pre-trained VAE discussed in Section 2.1. Also in LDM, the phoneme-level model significantly reduces computational demand compared to frame-level models by decreasing the computation required for iterative denoising, which is a major bottleneck of diffusion models. To simultaneously achieve naturalness and speaker similarity once by generating prior latent through LDM, our model is structured into two main components: conditioners and conditional diffusion model with dual CFG.

Conditioners include two types: the reference conditioner and the text conditioner. To inject conditional information for both the reference speaker and the text, these conditioners are designed to produce phoneme-wise conditions. These conditioners share inputs from the context encoder, which is a Transformer encoder employing cross-attention to model context-aware features derived from the outputs of the phoneme encoder and context embeddings. For obtaining context embeddings, we utilize pre-trained XLM-RoBERTa [21].

Given that the text conditioner relies solely on text inputs, we expect that by adjusting ω_{text} , the CFG weight for the text conditioner’s output c_{text} , we can control fine intelligibility. At the reference conditioner, we aim to generate speaker-aware phoneme-wise conditioning to facilitate zero-shot capability. To capture speaker’s style from reference speech and enable zero-shot capabilities, we integrate Retriever [22] into our reference conditioner. Reference speech is sampled from the target speaker’s subset, noise corrupted, and cut into random lengths to reduce training-inference mismatch. We extract NANSY features from the reference speech and then feed them into the cross-attention mechanism of the retriever encoder. The query for this cross-attention is fixed-length tokens, referred to as prototypes, which in our case is 60. Consequently, the output of this Transformer is also fixed-length tokens encapsulating the reference speech’s speaker style. Moreover, the reference conditioner encodes speaker-related conditions by leveraging these speaker tokens as a value for cross-attention and employs prototypes of identical length to those used by the retriever encoder. Analogous to ω_{text} in the text conditioner, the similarity to the speaker can be modulated by adjusting ω_{spk} , the CFG weight for reference conditioner’s output c_{spk} .

Our diffusion model’s architecture is also based on a Transformer encoder, akin to that of DiT [23]. Furthermore, instead of the conditioning mechanism from DiT’s adaptive layer norm, we change it to a simpler addition of conditions after two MLP layers, similar to that of DiffWave [24].

Our LDM is trained using the L_1 loss as WaveGrad [25]:

$$\mathcal{L} = \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mu + \sqrt{1 - \bar{\alpha}_t}\epsilon, t, c_{\text{spk}}, c_{\text{text}})\|_1, \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, I)$ is the added noise, ϵ_θ is the diffusion model with parameters, μ is the mean estimated by the VAE, t denotes the timestep, and $\bar{\alpha}_t$ corresponds to the noise coefficient at time t . We implement random dropout for both c_{text} and c_{spk} to employ CFG during the inference. Specifically, we drop c_{text} by 5% and c_{spk} by 10%, with an additional 10% dropout applied to both to promote the frequency of null-conditioned scenarios. Training employs a discrete integer diffusion timestep and a noise schedule. t is uniformly sampled from $[1, T]$, where $T = 200$. Following the approach of prior diffusion models [9, 24, 26], we adopt a linear variance schedule defined as $\beta_i = \beta_1 + (\beta_T - \beta_1)(i - 1)/(T - 1)$, setting $\beta_1 = 0.0001$

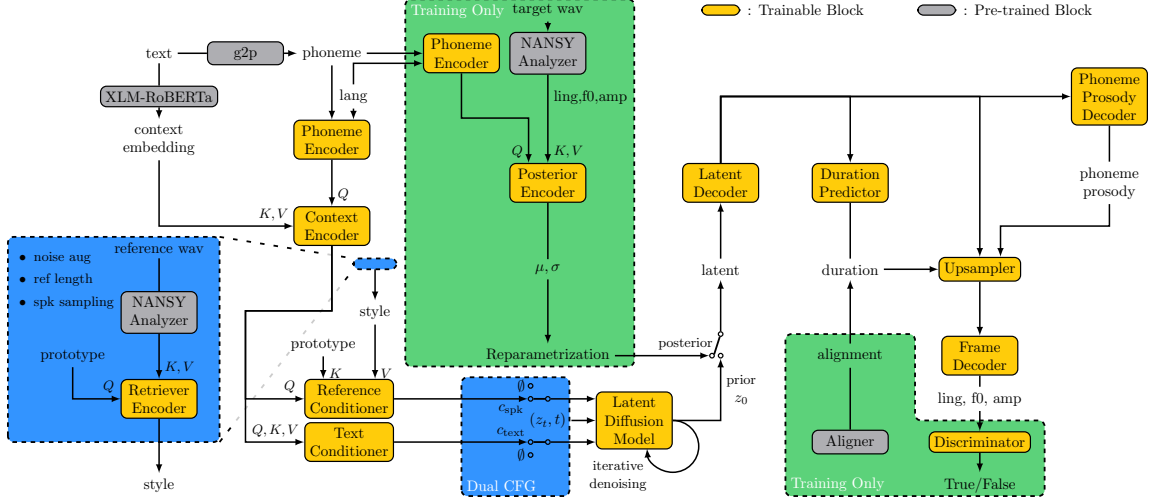


Figure 1: Overall model architecture of DualSpeech. Trainable blocks are colored in yellow and pre-trained modules are colored in gray. All blocks are based on the Transformer encoder architecture, even if their architecture is not mentioned in the main text.

and $\beta_T = 0.03$. The noise coefficient $\bar{\alpha}_t$ is calculated as $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i = \prod_{i=1}^t (1 - \beta_i)$.

2.3. Dual Classifier-Free Guidance for TTS

Our method extends to generate latent with fine control between text and reference conditions. Inspired by Lee *et al.* [11], DualSpeech employs a dual CFG for TTS, which can be represented as follows:

$$\begin{aligned} \tilde{\epsilon}_\theta(z_t, t, c_{\text{spk}}, c_{\text{text}}) &= \epsilon_\theta(z_t, t, c_{\text{spk}}, c_{\text{text}}) \\ &+ \omega_{\text{spk}} (\epsilon_\theta(z_t, t, c_{\text{spk}}, \emptyset) - \epsilon_\theta(z_t, t, \emptyset, \emptyset)) \\ &+ \omega_{\text{text}} (\epsilon_\theta(z_t, t, \emptyset, c_{\text{text}}) - \epsilon_\theta(z_t, t, \emptyset, \emptyset)), \end{aligned} \quad (2)$$

where $\tilde{\epsilon}_\theta$ represents the classifier free-guided noise, z_t is a latent at timestep t defined as $z_t = \sqrt{\bar{\alpha}_t} \mu + \sqrt{1 - \bar{\alpha}_t} \epsilon$, and \emptyset denotes zero tensors corresponding to a null-conditioned state. A noticeable difference from VoiceLDM lies in its conditioning from descriptions of the acoustic environment instead of the speaker’s style itself. This constraint is influenced by CLAP [27], which is trained not only with captions describing the speaker’s style but on a broad range of audio files and their captions. This includes the undesired noisy acoustic environments that are generally adverse to the objectives of TTS. In contrast to VoiceLDM, our approach allows for a more granular manipulation of speech synthesis, directly addressing the challenge of balancing text and speaker similarity.

2.4. Inference

During inference, only the LDM, the VAE decoder, and the NANSY synthesizer are utilized. The LDM generates phoneme-level latent through iterative denoising. We employ fast sampling suggested by Kong and Ping [26], utilizing a variance noise schedule as [1e-4, 5e-4, 1e-3, 5e-3, 0.01, 0.02, 0.05, 0.2, 0.3, 0.5, 0.4, 0.3, 0.3, 0.2, 0.1, 0.1]. The generated prior latent is then processed by the VAE decoder, which includes upsampling to the NANSY frame-level, and estimation of linguistic, f0, and amplitudes of NANSY. Finally, a raw waveform is synthesized by the pre-trained NANSY synthesizer.

3. Experiments

3.1. Settings

3.1.1. Training

All of our experiments were conducted on 8 NVIDIA RTX 4090 GPUs, utilizing dynamic batch sizes throughout the training process. Our pre-trained NANSY is trained with an identical setup following Choi *et al.* [7]. Our model processes input at a sampling rate of 16 kHz and generates outputs at a sampling rate of 44.1 kHz. The fundamental frequency is converted to the MIDI scale and then divided by 84, corresponding to 1,046 Hz. We applied an internal grapheme-to-phoneme (G2P) model to convert grapheme text to an IPA-based phoneme sequence.

3.1.2. Dataset

For the training of our TTS model, we utilized four datasets: LJSpeech [28], VCTK [29], Hi-Fi TTS [30], and LibriTTS [31]. These datasets encompass a wide range of speaker characteristics, including pronunciation, accents, timbre, and prosody, as well as linguistic nuances, offering a comprehensive diversity of English speech. The total dataset consists of 945 hours of high-quality speech and 2,576 English speakers, including a broad spectrum of English accents.

To evaluate how much the proposed method improves naturalness and similarity, a carefully curated set of 9 speakers exhibiting a broad spectrum of vocal characteristics was selected. The speaker set includes unseen non-English speakers and speakers with emotional tones such as sleepy.

Following previous studies [1, 2, 3, 4], we utilize a subset of the LibriSpeech test-clean dataset for objective evaluation. This subset contains speech clips with durations ranging from 4 to 10 seconds.

3.1.3. Evaluation Metrics

We assess the controllability of speaker similarity and naturalness through three distinct mean opinion scores (MOS): quality MOS (QMOS), which assesses sound quality, speaker similarity MOS (SMOS), which evaluates the similarity between the speaker of the prompt and the generated speech, and prosody MOS (PMOS), which gauges the naturalness of the speech’s prosody. To ensure a fair evaluation of audio with various sampling rates, all audio samples were downsampled to 16 kHz

Table 1: Subjective results for zero-shot TTS

	ω_{text}	ω_{spk}	QMOS	SMOS	PMOS
GT	-	-	3.92 \pm 0.15	-	4.34 \pm 0.13
YourTTS	-	-	2.41 \pm 0.13	1.40 \pm 0.12	1.98 \pm 0.14
HierSpeech++	-	-	3.62 \pm 0.12	3.31 \pm 0.19	3.01 \pm 0.16
StyleTTS 2	-	-	3.95 \pm 0.12	1.84 \pm 0.17	3.76 \pm 0.14
DualSpeech (Ours)	1.0	4.0	4.18 \pm 0.11	3.74 \pm 0.21	3.36 \pm 0.17
DualSpeech (Ours)	4.0	1.0	4.24 \pm 0.12	2.35 \pm 0.22	3.83 \pm 0.15

before being assessed. For evaluation, the LibriTTS test-other subset was utilized as the input text. To assess the ground truth (GT) for QMOS, reference speeches that were used as speaker prompts were evaluated. Similarly, for PMOS GT evaluation, LibriTTS speech samples corresponding to the input text were measured.

To assess the correctness and intelligibility of the generated speech, we measure the word error rate (WER) and character error rate (CER) by comparing the transcribed text of the generated speech with the corresponding input text. We transcribe speech by pre-trained CTC-based HuBERT-Large² [32].

3.2. Results

3.2.1. Subjective Evaluation

We conducted three MOS tests for subjective evaluation, comparing our model against models with official implementations, including those that are state-of-the-art models [5, 6, 33]. The distinctive feature of DualSpeech lies in its capability to precisely modulate the balance between text content and speaker characteristics with the CFG weights (ω_{text} and ω_{spk}), thereby enabling synthesis to adapt to diverse application scenarios.

Our system consistently maintains a high level of quality in terms of QMOS, while also demonstrating the ability to selectively enhance either PMOS or SMOS through strategic weight adjustments. In configurations prioritizing text content with $(\omega_{\text{text}}, \omega_{\text{spk}}) = (4.0, 1.0)$, our system not only achieves a QMOS of 4.24, indicative of superior sound quality but also achieves a PMOS of 3.83, underscoring its exceptional proficiency in replicating natural prosody. This suggests our system’s adeptness at capturing and reproducing the nuanced tones and rhythms inherent in phonemes. Moreover, it enables us to faithfully replicate the timbre of speakers while excluding biased expressions found in reference speech, such as yawning, and instead generate neutral expressions derived from the training datasets with the reference’s timbre.

Conversely, when emphasizing speaker characteristics with $(\omega_{\text{text}}, \omega_{\text{spk}}) = (1.0, 4.0)$, while maintaining the same level of QMOS, our system significantly improves SMOS to 3.74. This underscores its outstanding ability to capture speaker similarity. The noticeable enhancement in SMOS accentuates the system’s capability to replicate distinct voice traits of speakers, which is vital for personalized voice synthesis applications.

3.2.2. Objective Result

In TTS studies, objective measures like WER and CER serve as critical benchmarks for evaluating the robustness and correctness of synthesized speech. Our study presents an extensive objective evaluation conducted on a subset of the LibriSpeech dataset, the results of which are detailed in Table 2. This evaluation underscores the efficacy of our proposed TTS system, particularly when compared against state-of-the-art systems [1, 2, 3, 4, 33] and GT.

²<https://huggingface.co/facebook/hubert-large-ls960-ft>

Table 2: Evaluating LibriSpeech-subset for robustness and accuracy with highest scores in bold, second highest underlined, and baseline scores asterisked.

	ω_{text}	ω_{spk}	WER \downarrow	CER \downarrow
GT	-	-	2.26	0.61
YourTTS* [33, 3]	-	-	7.92	3.18
VALL-E* [1]	-	-	5.9	-
Voicebox* [2]	-	-	1.90	-
CLaM-en* [3]	-	-	5.11	2.87
SPEAR-TTS* [4]	-	-	-	1.92
DualSpeech (Ours)	1.0	1.0	2.77	0.83
DualSpeech (Ours)	1.0	2.0	2.62	0.81
DualSpeech (Ours)	2.0	1.0	<u>2.59</u>	0.77
DualSpeech (Ours)	2.0	2.0	2.62	<u>0.80</u>

Table 3: Inference speed of models. CLaM-en and DualSpeech were tested on an A100, and Voicebox’s GPU details are undisclosed.

	Voicebox* [2]	CLaM-en* [3]	DualSpeech (Ours)
Inference Time (s)	6.4 (64 NFE)	4.2	0.19 (16 NFE)

The GT recordings exhibit low WER and CER, at 2.26 and 0.61 respectively, setting a high standard for speech synthesis quality. Among the competing systems, Voicebox achieves an impressive WER of 2.00 in one instance, the lowest among the synthesized voices, albeit without a corresponding CER reported. Our system, under various configurations of text and speaker CFG weights, demonstrates competitive performance, particularly with a configuration of $(\omega_{\text{text}}, \omega_{\text{spk}}) = (2.0, 1.0)$, achieving a WER of 2.59 and a CER of 0.77. These results are notably close to the GT, highlighting our system’s ability to maintain high levels of speech intelligibility and accuracy.

Furthermore, our system’s adaptability is evident in its performance across different configurations, suggesting that precisely adjusting the balance between text and speaker emphasis can optimize performance for specific applications. While no single configuration universally outperforms all others, the ability to adjust these parameters allows for significant flexibility in tailoring the system to meet diverse needs.

In addition to speech synthesis quality, inference speed is a crucial factor for the practical application of TTS systems. Table 3 shows our phoneme-level diffusion model’s superior inference speed, clocking in at 0.19 seconds, significantly faster than other frame-based diffusion models or auto-regressive language models, including Voicebox and CLaM-en, which report inference times of 6.4 and 4.2 seconds, respectively. This remarkable speed does not compromise the quality of the synthesized speech, positioning our system as a highly efficient and effective solution for real-time TTS applications.

4. Conclusion

In this work, we introduce DualSpeech, a text-to-speech model that combines a phoneme-level latent diffusion model with dual classifier-free guidance (CFG). This model showcases exceptional zero-shot TTS capabilities, excelling in speaker-fidelity and text-intelligibility. DualSpeech provides high-quality voice synthesis with the flexibility to adjust for either speaker-fidelity or text-intelligibility, according to specific requirements. We believe that integrating our dual CFG approach into any diffusion-based TTS system will significantly refine the balance between speaker fidelity and text intelligibility.

5. References

- [1] C. Wang, S. Chen, Y. Wu, Z.-H. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Neural codec language models are zero-shot text to speech synthesizers," *ArXiv*, vol. abs/2301.02111, 2023.
- [2] M. Le, A. Vyas, B. Shi, B. Karrer, L. Sari, R. Moritz, M. Williamson, V. Manohar, Y. Adi, J. Mahadeokar, and W.-N. Hsu, "Voicebox: Text-guided multilingual universal speech generation at scale," in *NeurIPS*, vol. 36, 2023.
- [3] J. Kim, K. Lee, S. Chung, and J. Cho, "CLaM-TTS: Improving neural codec language model for zero-shot text-to-speech," in *ICLR*, 2024.
- [4] E. Kharitonov, D. Vincent, Z. Borsos, R. Marinier, S. Girgin, O. Pietquin, M. Sharifi, M. Tagliasacchi, and N. Zeghidour, "Speak, read and prompt: High-fidelity text-to-speech with minimal supervision," *TACL*, vol. 11, pp. 1703–1718, 2023.
- [5] S.-H. Lee, H.-Y. Choi, S.-B. Kim, and S.-W. Lee, "Hierspeech++: Bridging the gap between semantic and acoustic representation of speech by hierarchical variational inference for zero-shot speech synthesis," *arXiv preprint arXiv:2311.12454*, 2023.
- [6] Y. A. Li, C. Han, V. S. Raghavan, G. Mischler, and N. Mesgarani, "StyleTTS 2: Towards human-level text-to-speech through style diffusion and adversarial training with large speech language models," in *NeurIPS*, vol. 36, 2023.
- [7] H.-S. Choi, J. Yang, J. Lee, and H. Kim, "NANSY++: Unified voice synthesis with neural analysis and synthesis," in *ICLR*, 2023.
- [8] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, "Neural analysis and synthesis: Reconstructing speech from self-supervised representations," in *NeurIPS*, vol. 34, 2021, pp. 16 251–16 265.
- [9] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," in *NeurIPS*, 2020, pp. 6840–6851.
- [10] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *NeurIPS Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [11] Y. Lee, I. Yeon, J. Nam, and J. S. Chung, "VoiceLDM: Text-to-Speech with Environmental Context," *arXiv preprint arXiv:2309.13664*, 2023.
- [12] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR*, 2014.
- [13] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *IEEE/CVF CVPR*, 2022, pp. 10 684–10 695.
- [14] J. Kim, S. Kim, J. Kong, and S. Yoon, "Glow-TTS: A Generative Flow for Text-to-Speech via Monotonic Alignment Search," in *NeurIPS*, vol. 33, 2020, pp. 8067–8077.
- [15] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, A. Rosenberg, B. Ramabhadran, and Y. Wu, "Generating diverse and natural text-to-speech samples using a quantized fine-grained vae and autoregressive prosody prior," in *IEEE ICASSP*, 2020, pp. 6699–6703.
- [16] Z. Liu, Y. Guo, and K. Yu, "Diffvoice: Text-to-speech with latent diffusion," in *IEEE ICASSP*, 2023, pp. 1–5.
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, vol. 30, 2017.
- [18] I. Elias, H. Zen, J. Shen, Y. Zhang, Y. Jia, R. Skerry-Ryan, and Y. Wu, "Parallel Tacotron 2: A Non-Autoregressive Neural TTS Model with Differentiable Duration Modeling," in *Proc. Interspeech*, 2021, pp. 141–145.
- [19] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, vol. 27, 2014, pp. 2672–2680.
- [20] J. Bae, J. Yang, T. Bak, and Y.-S. Joo, "Hierarchical and multi-scale variational autoencoder for diverse and natural non-autoregressive text-to-speech," in *Proc. Interspeech*, 2022, pp. 813–817.
- [21] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov, "Unsupervised cross-lingual representation learning at scale," in *ACL*, 2019.
- [22] D. Yin, X. Ren, C. Luo, Y. Wang, Z. Xiong, and W. Zeng, "Retriever: Learning content-style representation as a token-level bipartite graph," in *ICLR*, 2022.
- [23] W. Peebles and S. Xie, "Scalable diffusion models with transformers," in *IEEE/CVF ICCV*, October 2023, pp. 4195–4205.
- [24] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "Diffwave: A versatile diffusion model for audio synthesis," in *ICLR*, 2021.
- [25] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "Wavegrad: Estimating gradients for waveform generation," in *ICLR*, 2021.
- [26] Z. Kong and W. Ping, "On fast sampling of diffusion probabilistic models," in *ICML Workshop on INNF*, 2021.
- [27] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, "Clap: Learning audio concepts from natural language supervision," in *IEEE ICASSP*, 2023, pp. 1–5.
- [28] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [29] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "Superseded-ctr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (Version 0.92)," 2016. [Online]. Available: <https://datashare.ed.ac.uk/handle/10283/3443>
- [30] E. Bakhturina, V. Lavrukhin, B. Ginsburg, and Y. Zhang, "Hi-Fi Multi-Speaker English TTS Dataset," in *Proc. Interspeech*, 2021, pp. 2776–2780.
- [31] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech," in *Proc. Interspeech*, 2019, pp. 1526–1530.
- [32] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM TASLP*, vol. 29, p. 3451–3460, oct 2021.
- [33] E. Casanova, J. Weber, C. D. Shulby, A. C. Junior, E. Gölge, and M. A. Ponti, "YourTTS: Towards zero-shot multi-speaker TTS and zero-shot voice conversion for everyone," in *ICML*, 2022, pp. 2709–2720.